An Analysis of the Linear Bilateral ANAC Domains Using the MiCRO Benchmark Strategy

Dave de Jonge
IIIA-CSIC, Barcelona, Spain
davedejonge@iiia.csic.es

Abstract

The Automated Negotiating Agents Competition (ANAC) is an annual competition that compares the state-of-the-art algorithms in the field of automated negotiation. Although in recent years ANAC has given more and more attention to more complex scenarios, the linear and bilateral negotiation domains that were used for its first few editions are still widely used as the default benchmark in automated negotiations research. In this paper, however, we argue that these domains should no longer be used, because they are too simplistic. We demonstrate this with an extremely simple new negotiation strategy called MiCRO, which does not employ any form of opponent modeling or machine learning, but nevertheless outperforms the strongest participants of ANAC 2012, 2013, 2018 and 2019. Furthermore, we provide a theoretical analysis which explains why MiCRO performs so well in the ANAC domains. This analysis may help researchers to design more challenging negotiation domains in the future.

1 Introduction

The field of automated negotiation deals with autonomous agents that are purely self-interested, but still need to cooperate to ensure beneficial outcomes. Each agent may propose potential solutions to the other agents, which may then accept or reject those proposals. Although each agent is purely self-interested, its proposals must also benefit the others because otherwise they would never accept them. Therefore, a negotiating agent must strike a balance between maximizing its own utility and the utility of its opponents [Faratin *et al.*, 1998]. One of the main questions studied in automated negotiation is how to make this trade-off. That is, how to find the optimal *concession strategy*. The simplest types of concession strategy start by making very selfish proposals, but, as time passes, concede by making proposals that yield more and more utility to the opponent. Many variants of this basic strategy have been proposed, which often require knowledge about the opponent's strategy and utility function. Although such knowledge is typically not given, the agent may try to infer it, at run-time, from the proposals it receives from its opponent. This is known as *opponent modeling*. Therefore, a second main question studied in automated negotiation, is how to implement effective opponent modeling algorithms.

To tackle these questions, the annual Automated Negotiating Agent Competition (ANAC) has been organized since 2010. The first three editions of this competition focused on basic bilateral negotiations with linear utility functions [Baarslag *et al.*, 2012], and in 2013 the option was added for agents to learn from previous negotiation sessions [Gal and Ilany, 2015]. In 2014 the focus shifted to very large domains with non-linear utility functions [Fujita *et al.*, 2014]. However, from 2015 onward the competition went back to smaller domains and linear utility, but focused on multilateral negotiations [Fujita *et al.*, 2017]. In 2019 the focus shifted back to bilateral negotiations, but this time with only *partially* known utility functions [Aydogan *et al.*, 2020]. Furthermore, since 2017 ANAC has been extended with a number of separate leagues focused on more specific challenges, such as the game of Diplomacy [Jonge *et al.*, 2019], supply chain environments [Aydogan *et al.*, 2020], and the game of Werewolves [Aydogan *et al.*, 2020].

Despite the increased attention to more complex scenarios, the most basic linear negotiation domains from ANAC 2010-2013 are still widely used as the default benchmark for automated negotiation. For example, [Sengupta *et al.*, 2021] used the ANAC 2013 domains to test their work, while [Mirzayi *et al.*, 2021] used a small selection of domains from both ANAC 2012 and 2013, and [Bakker *et al.*, 2019] used similar linear ANAC-style domains. Furthermore, the domains used for the multilateral and partial preference leagues of ANAC 2018 and 2019 were also still based on such linear domains.

Even for the basic settings of ANAC 2010-2013, a plethora of different opponent modeling techniques have been applied by the participants, such as Gaussian processes, Bayesian learning, reinforcement learning, wavelet decomposition and cubic smoothing splines. A survey of such techniques can be found in [Baarslag *et al.*, 2016].

Despite the fact that opponent-modeling algorithms have proved very successful in the ANAC competitions, we argue that these results may not be as meaningful as previously thought, because the domains that were used in ANAC can be tackled just as well without opponent modeling. To demonstrate this, we present a new algorithm, called MiCRO, which does not employ any kind of opponent modeling or machine learning. We show experimentally that, despite its simplicity, it outperforms the strongest participants of ANAC. Furthermore, we provide a theoretical analysis that shows that in a certain class of domains MiCRO is a near-optimal strategy, and that many domains used in ANAC belong to this class.

We should stress, however, that we are certainly not claiming that opponent modeling is not important at all. On the contrary, we still think that opponent modeling is an essential aspect of automated negotiation. But we do argue that any conclusions drawn in the past solely from experiments with linear ANAC domains may need to be reviewed, and that in the future any new algorithms should either be tested on more complex test cases, such as the non-linear ANAC 2014 domains or the Diplomacy game, or otherwise should at least be compared with MiCRO. Furthermore, our conclusions do not only apply to opponent modeling. There are also other aspects of negotiation that are typically applied by negotiation algorithms, but not by MiCRO. For example, MiCRO does not make use of the fact that the utility functions are linear. In fact, it does not even need to know its own utility function at all. It only needs to know its preference ordering over the domain's offers.

The source code of MiCRO has been made publicly available at: https://www.iiia.csic.es/~davedejonge/downloads

2 The Theory of Negotiation

In a classical scenario for automated negotiation two agents a_1 and a_2 are bargaining to agree on some contract. The agents have a fixed amount of time to make proposals to one another. That is, each agent may propose an offer ω, from some given set of possible offers Ω, to the other agent which may then either accept the proposal or reject it and make a counter proposal $\omega' \in \Omega$. The agents continue making proposals to each other until either the deadline has passed, or one of the agents accepts a proposal made by the other. Each agent a_i has a utility function u_i that assigns to each offer $\omega \in \Omega$ a utility value $u_i(\omega) \in \mathbb{R}$, but which is not known to the other agent. When an offer ω gets accepted the agents receive their respective utility values $u_1(\omega)$ and $u_2(\omega)$ corresponding to this offer. On the other hand, if the negotiations fail because no proposal was accepted before the deadline, then each agent a_i receives a fixed utility value $rv_i \in \mathbb{R}$, which is known as its *reservation value*.

In some cases, the negotiation domain is also assumed to come with a so-called *discount factor*. In that case, the utility obtained by the agents also depends on the time at which the agreement was made. The later the agreement is made, the lower the utility. That is: $\hat{u}_i(\omega, t) = u_i(\omega) \cdot \delta_i^t$ where $\hat{u}_i(\omega, t)$ is the actual utility obtained by agent a_i if the agents come to an agreement ω at time $t \in (0, 1]$, and where $\delta_i \in (0, 1]$ is the discount factor of a_i.

Definition 1. *A bilateral* ***negotiation domain*** *is a tuple* $\langle \Omega, u_1, u_2, rv_1, rv_2, \delta_1, \delta_2 \rangle$ *where:*

- Ω *is the set of possible* ***offers****.*
- u_1 *and* u_2 *are two undiscounted* ***utility functions*** *(one for each agent) which are maps from* Ω *to* $[0, 1]$
- rv_1*,* $rv_2 \in [0, 1)$ *are the* ***reservation values*** *of the respective agents.*
- δ_1*,* $\delta_2 \in (0, 1]$ *are the* ***discount factors*** *of the respective agents.*

2.1 Time-based Strategies

The simplest types of negotiating agents follow a fixed time-based concession strategy and do not adapt to their opponents [Faratin *et al.*, 1998]. Time-based agents apply a so-called *aspiration function* $f(t)$ that typically decreases over time. Whenever it is the agent's turn to make a proposal, the agent will propose an offer ω for which its utility $u_i(\omega)$ is equal (or close to) $f(t)$. We can distinguish between two types of such agents: hard-headed ones, and conceding ones. Hard-headed agents concede very little, so as to maximize their profit when a deal is made, but at the risk of failing to make a deal at all. Conceding agents, on the other hand, are willing to make larger concessions, to increase the likelihood of making a deal, but at the price of making less profit from those deals.

2.2 Adaptive Strategies

Most agents that have been successful in ANAC, however, take a more intelligent approach: they assume the opponent plays a time-based strategy, and then try to predict how far the opponent is willing to concede. The adaptive agent then makes sure it will never propose or accept any offers with utility lower than the maximum utility the opponent is predicted to offer.

If the opponent is indeed time-based and the prediction is accurate, then the adaptive agent will obtain the maximum score it could theoretically get against that opponent. However, it should be noted that a *purely* adaptive strategy is not game-theoretically stable. If all opponents in a tournament are time-based, then an adaptive strategy is the best response, but if all opponents use an adaptive strategy, then a hard-headed time-based strategy becomes the best response because all opponents would concede to its strong demands. Therefore, the question which strategy is the best, is a kind of chicken-and-egg problem.

One solution to this, is to use a strategy that combines being hard-headed and being adaptive. The idea is that the agent a_i chooses a certain threshold value $\theta_i \in \mathbb{R}$, and although the agent adapts to its opponent, it will not propose or accept any offer ω for which its utility $u_i(\omega)$ is below this threshold. So, θ_i is the lowest value the agent is willing to concede to. Of course, this then raises the question how the agent could determine the optimal value for θ_i, but this is difficult to answer since it depends on the opponent's utility function, which is not (perfectly) known.

2.3 Tit-for-Tat Strategies

Another approach is the so-called Tit-for-Tat (TFT) negotiation strategy [Faratin *et al.*, 1998]. Instead of adapting to the opponent, TFT mimics the opponent. Every time the opponent concedes, TFT tries to make an concession of equal size. The idea is that rather than assuming the opponent uses a fixed time-based strategy, it assumes the opponent is able to adapt, so TFT can try to force the opponent to concede more. The disadvantage of TFT is that if the opponent concedes too much, then TFT cannot really exploit this, but the advantage is that it cannot really *be exploited* either.

2.4 The MiCRO Strategy

We now introduce an entirely new concession strategy, which we call MiCRO, which stands for *Minimal Concession in Reply to new Offers*. Simply stated, it works as follows: whenever the opponent proposes a *new* offer, MiCRO also replies with a new offer. This offer will be the one with highest utility for the agent itself, that the agent has not yet proposed before. On the other hand, when the the opponent repeats an offer it has already proposed before, then MiCRO also replies with an offer it has already proposed before.

More formally, let a_1 denote an agent that applies the MiCRO strategy, and a_2 its opponent (which may be applying any arbitrary strategy), and let $K := |\Omega|$ denote the size of the domain. Before the negotiations begin, our agent a_1 creates a list $(\omega_1, \omega_2, \dots \omega_K)$ containing all offers in the domain, sorted in order of decreasing utility for itself. That is, $u_1(\omega_1) \geq u_1(\omega_2) \geq \dots \geq u_1(\omega_K)$. Then, whenever it is a_1's turn to make a proposal, it counts how many *different* offers it has so far received from the opponent (we denote this number by n), and how many *different* offers it has so far proposed to the opponent (we denote this number by m). If $m \leq n$ then it proposes ω_{m+1}. On the other hand, if $m > n$ then it picks a random integer r such that $1 \leq r \leq m$ and proposes ω_r. Of course, it should never propose any offer that is below its reservation value, so in case $u_1(\omega_{m+1}) < rv_1$, it also just repeats a random previous proposal, even if $m \leq n$. The important thing to note, is that *MiCRO does not apply any form of opponent modeling*, and, in fact, it does not even need to know its own utility function. It only needs to have a full preference ordering over its offers.

Although MiCRO is somewhat similar to TFT in the sense that it mimics the behavior of the opponent, there are several major differences. The first difference is that MiCRO does not care *how large* the opponent's concessions are. The motivation for this is that we assume the opponent's utility is unknown, and therefore the size of the opponent's concession as perceived by our agent says nothing about the size of the concession the opponent *intended* to make. The opponent might make a large concession in terms of its own utility, but this may result in a very small concession measured in our agent's utility. For the same reason, MiCRO never makes large concessions to its opponent. In fact, it always makes exactly the smallest possible concession: it just proposes the next offer on its list. The second difference between MiCRO and TFT, is that MiCRO uses a different definition of 'concession'. That is, even if a_2's new proposal offers less utility to a_1 than a_2's previous proposal, MiCRO still considers this a concession, as long as the offer is *different* from any of a_2's previous offers. Again, this is because it may still have been a concession in terms of the a_2's utility.

2.5 Acceptance Strategy

MiCRO can be combined with various acceptance strategies, but in this paper we will assume it is always combined with a simple acceptance strategy that accepts a received offer if and only if it is better than or equal to the lowest offer it is, at that time, willing to propose.

More precisely, if agent a_1 applies MiCRO and we define:

$$\omega_{low} := \begin{cases} \omega_{m+1} & \text{if } m \leq n \\ \omega_m & \text{if } m > n \end{cases} \tag{1}$$

(with m and n defined as before) then a received offer ω is accepted by a_1 iff $u_1(\omega) \geq \max\{u_1(\omega_{low}), rv_1\}$.

3 Measuring Performance

Before we present our experiments, we here discuss the two different ways in which we evaluated the agents in our experiments, namely by means of a *tournament evaluation* and by means of a *game-theoretical evaluation.*

Both methods involve running a large number of negotiation sessions in which each agent negotiates against every other agent, in a number of different domains.

The tournament evaluation simply consist in calculating the average utility obtained by each agent over all sessions in which it participated (one may or may not choose to include the sessions in which an agent played against itself). The main downside of tournament evaluation, is that every result against every opponent counts the same. This is arguably somewhat unrealistic because in a real-world situation one is more likely to encounter stronger opponents than to encounter weaker opponents. After all, if one strategy is known to be weak, it is less likely that anyone would ever employ that strategy. For this reason, several authors have argued in favor of a game-theoretical evaluation [Williams *et al.*, 2011; Baarslag *et al.*, 2013; Chen and Weiss, 2013].

For the game-theoretical evaluation we create a matrix that contains the average scores that each agent obtained against each opponent separately. This matrix can then be seen as the pay-off matrix of a symmetric normal-form game in which the two players each choose one of the agents as their strategy. We can then determine if any choice of agents constitutes a Nash equilibrium. The downside, however, is that a pure Nash equilibrium may not always exist.

We calculated the scores of the agents in our experiments in exactly the same way as in the ANAC tournaments. To explain how this calculation works, we define the notion of an ***encounter*** as a tuple (a_i, a_j, d, r) where a_i and a_j are two agents, d is a negotiation domain, and r is an integer that acts as an identifier to distinguish different negotiation sessions with the same agents and domain. So, (a_i, a_j, d, r) is the r^{th} encounter in which agent a_i and a_j negotiate against each other in domain d, with a_i having utility function u_1 and a_j having utility function u_2 of that domain. Note that the tuples (a_i, a_j, d, r) and (a_j, a_i, d, r) represent different encounters, with the difference being that the two utility functions are switched between the two agents. The utility value $u_k(a_i, a_j, d, r)$ obtained by a_k in encounter (a_i, a_j, d, r), with either $a_k = a_i$, or $a_k = a_j$, is defined as:

$$u_k(a_i, a_j, d, r) := \begin{cases} u_k(\omega) \cdot \delta_k^t & \text{if an agreement is made} \\ rv_k & \text{if no agreement is made} \end{cases}$$

where δ_k and rv_k are the respective discount factor and reservation value of agent a_k in domain d, ω is the offer the two agents agree upon, and t is the time at which they make the agreement (in that specific encounter).

Then, we calculate the score $u_k(a_i)$ that an agent a_k obtains against opponent a_i (which may or may not be the same agent) over an entire experiment as:

$$u_k(a_i) := \frac{1}{2|D|} \sum_{d \in D} \frac{1}{R_{k,i,d}} \quad \times \\ \Big(\sum_{r=1}^{R_{k,i,d}} u_k(a_k, a_i, d, r) + \sum_{r=1}^{R_{k,i,d}} (u_k(a_i, a_k, d, r) \Big) \tag{2}$$

where D is the set of negotiation domains and $R_{k,i,d} \in \mathbb{N}$ is the number of times the negotiation is repeated with agents a_k, a_i and domain d. The first summation between the parentheses represents all encounters in which a_k has utility u_1 and the second summation represents all encounters in which a_k has utility u_2.

4 Empirical Evidence

In this section we present the empirical evidence for our claim that the ANAC domains are too simple. Specifically, we present our experiments that show that MiCRO outperforms the best agents from several editions of ANAC. All experiments were conducted using the Genius platform [Lin *et al.*, 2014], on a laptop with Intel Core i7-8750H@2.20GHz CPU and 32 GB RAM. In all cases the deadlines were set to 180 seconds (which is the standard in ANAC). For each domain and each pair of agents the negotiations were repeated at least 3 times, and after that, if necessary, repeated several times more to ensure the standard error on all the values displayed in the tables was always lower than 0.01. In our tables all results with standard error lower than 0.001 are displayed with 3 digits after the decimal point, while results with standard error between 0.001 and 0.01 are displayed with 2 decimals.

4.1 ANAC 2018

In our first experiment we ran a tournament between the top-3 agents of ANAC 2018 and MiCRO. Since ANAC 2018 involved multilateral domains while we are focusing only on bilateral negotiations, we used the domains of ANAC 2013 in this experiment (all 18 of them).

The results are displayed in Table 1. Each entry in the first four columns shows the score $u_k(a_i)$ as defined by Eq. (2), with a_k the agent indicated in the row header, and its opponent a_i indicated in the column header. For example, in the sessions in which MiCRO negotiated against BetaOne, MiCRO scored on average 0.685 points (the intersection of the row labeled MiCRO and the column labeled BetaOne), while BetaOne scored 0.686 (the row labeled BetaOne and the column labeled MiCRO). The last two columns show the average score of each agent against all opponents, respectively with and without the score of the agent when playing against itself taken into account (calculated as $\frac{1}{4}\sum_{a_i} u_k(a_i)$ and $\frac{1}{3}\sum_{a_i, i\neq k} u_k(a_i)$ respectively). In each column the best result is emphasized in bold.

We see that MiCRO scores the highest tournament score, both with and without self-play. Furthermore, if we consider this table as the pay-off matrix of a game, then MiCRO vs. MiCRO forms a pure Nash equilibrium of that game. We can see this as follows. Suppose the opponent is AgreeableAgent2018. Then the best response to this opponent can be found by looking at that opponent's column (labeled Agr.Ag'18), and finding the highest score in this column. We see the highest score in this column is obtained by MiCRO (0.583 points), so MiCRO is the best response against AgreeableAgent2018. In this way we can also see that MiCRO is the best response against itself, so MiCRO vs. MiCRO is indeed a Nash equilibrium. Note that Beta One vs. Beta One is also a Nash Equilibrium, but that MiCRO vs. MiCRO is better, because both agents receive 0.745 points, while in the other equilibrium both players only receive 0.708 points.

	MiCRO	B. One	Agr.Ag.'18	M. Wan	Av.	Av.ex.s.p.
MiCRO	**0.745**	0.685	**0.583**	0.65	**0.667**	**0.641**
Beta One	0.686	**0.708**	0.532	**0.69**	0.654	0.636
Agr.Ag.'18	0.56	0.621	0.57	0.64	0.598	0.607
Meng Wan	0.54	0.64	0.58	0.54	0.57	0.585

Table 1: Results of our mini tournament with the top agents of ANAC 2018. 'Av.' stands for Average Score, and 'Av. ex. s.p.' stands for Average Score Excluding Self-Play.

To find the *mixed* symmetric Nash equilibria we used the Gambit library [McKelvey *et al.*, 2017]. We found one such equilibrium, in which each player chooses MiCRO with 28% probability and Beta One with 72% probability. This yields an expected value of 0.702 for each player, so it is even worse than the two pure equilibria.

4.2 ANAC 2019

Next, we ran a tournament with the top-3 agents of ANAC 2019. The challenge of ANAC 2019 was to negotiate with only a partial preference ordering over the possible deals, rather than a utility function. However, we are only interested in comparing concession strategies and the the question how to deal with partial preferences is beyond the scope of our work. Therefore, we took an intermediate approach: we provided the agents with *full* preference orderings, rather than partial ones, so we could still use MiCRO. We again used the domains of ANAC 2013, but this time with all reservation values set to 0 and all discount factors[1] to 1. Unfortunately, the agent KakeSoba crashed on all domains with more than 1000 offers, so we could only use the 12 smallest domains.

The results are displayed in Table 2. In this case, MiCRO ends in a shared first place with AgentGG in the tournament evaluation, both with and without self-play (MiCRO scores slightly higher if we round off to 3 decimals, but the difference is not statistically significant). Furthermore, we note that MiCRO vs. MiCRO forms the only pure Nash equilibrium, and we checked with Gambit that there are no mixed symmetric Nash equilibria.

[1]For this experiment we set the discount factors to 1 because ANAC 2019 itself did not involve discount factors, and we set the reservation values to 0 because otherwise MiCRO would not be able to know which offers were rational. For all other experiments we kept the original reservation values and discount factors.

	MiCRO	Agent GG	K.Soba	SAGA	Av.	Av.ex.s.p.
MiCRO	**0.753**	**0.78**	0.48	**1.000**	**0.75**	**0.75**
Agent GG	0.70	0.747	**0.54**	**1.000**	**0.75**	**0.75**
KakeSoba	0.45	0.65	0.44	**1.00**	0.63	0.70
SAGA	0.427	0.420	0.43	0.70	0.49	0.42

Table 2: Results of our experiment with the top agents of ANAC 2019.

	MiCRO	Ag.LG	OMAC	CUHK	T.N.R	Av.	Av.ex.s.p.
MiCRO	**0.790**	**0.668**	0.629	0.611	0.745	**0.689**	**0.663**
AgentLG	0.61	0.64	0.64	**0.63**	0.757	0.66	**0.66**
OMAC	0.591	0.62	0.63	0.58	0.737	0.63	0.631
CUHK	0.508	0.62	**0.66**	0.59	**0.77**	0.63	0.64
T.N.R.	0.643	0.61	0.59	0.50	0.69	0.61	0.58

Table 3: Results of our experiment with the top agents of ANAC 2012.

4.3 ANAC 2012, ANAC 2013, and Boulware

Above, we have compared MiCRO with the most recent agents that are available with the Genius framework. However, we should note that these agents were actually implemented for slightly different settings. After all, ANAC 2019 involved negotiations with partial preference orderings, and ANAC 2018 involved multilateral negotiations. Although this is not an uncommon approach (e.g. see [Sengupta *et al.*, 2021; Mirzayi *et al.*, 2021]), one could argue that this is not the best way to compare MiCRO. For this reason, we have also tested MiCRO against the top agents of ANAC 2012 and 2013, which were the last two editions which involved the simplest bilateral negotiation scenario.

For our experiment with the ANAC 2012 agents we used the top-4 of that year (because there was no statistically significant difference between the number 3 and number 4 of that year's competition). In [Williams *et al.*, 2014] the domains of ANAC 2012 were classified as having either *low*, *medium* or *high* competitiveness. For our experiments we used the three largest domains of each of these three classes, so nine domains in total.[2] The results are displayed in Table 3. In this case MiCRO wins the tournament with self-play, and ends in a shared first place with AgentLG if we exclude self-play. Furthermore, MiCRO vs. MiCRO is the only pure Nash Equilibrium.

For our experiment with ANAC 2013 we used the top-3 agents of that year, and all 18 domains. The results are displayed in Table 4. Once again, MiCRO ends in first place, with or without self-play (altough the differences are small), and MiCRO vs. MiCRO is the only pure Nash equilibrium.

For both experiments we also checked with Gambit that there are no mixed symmetric equilibria.

Finally, we tested MiCRO against the two implementations of a Boulware strategy that are included in Genius. We again used all domains of ANAC 2013 and found that MiCRO outperformed both agents, with or without self-play.

[2]In ANAC 2012 each domain had three variants, with different reservation values and discount factors. These variants were labeled with A, B and C respectively. For our experiments we used the variants labeled with A.

	MiCRO	MetaAg.	TMFAg.	TheFawkes	Av.	Av.ex.s.p.
MiCRO	**0.745**	**0.71**	**0.60**	**0.69**	**0.68**	**0.66**
MetaAgent	0.65	0.68	**0.60**	0.68	0.65	0.64
TMFAgent	0.47	0.67	0.58	0.64	0.59	0.59
TheFawkes	0.56	0.67	0.53	0.395	0.54	0.59

Table 4: Results of our experiment with the top agents of ANAC 2013.

4.4 Speed

One might expect MiCRO to be very slow because it only makes minimal concessions. But, on the other hand, MiCRO saves a lot of time because it does not have to update any opponent models. When we look at negotiation sessions in which both sides were played by the same agent, it turns out that those that involved MiCRO generally finished much faster than the others. In fact, even in the largest domain in our experiments, the *Energy* domain with 390,625 offers, the MiCRO vs. MiCRO sessions always ended in agreement, and always within 80 seconds or less.

5 Theoretical Evidence

In this section we present the theoretical evidence for our claim that the ANAC domains are too simple. We do this in three steps. First, we define for any negotiation domain a pair of values that we call the *balance values*. Next, we show that these are the lowest values to which MiCRO is willing to concede. And finally, we show that for most ANAC domains the balance values lie very close to the Nash bargaining solution. Together, this means that MiCRO plays a near-optimal strategy in such domains. Due to space restrictions we cannot provide the proofs of our propositions, but we will publish them later in an extended version of this paper.

In the following, for any real number x we define Ω_i^x to be the set of all offers for which $u_i(\omega) \geq x$.

$$\Omega_i^x := \{\omega \in \Omega \mid u_i(\omega) \geq x\}$$

Note that the set of offers that MiCRO is willing to propose or accept is always of the form Ω_i^x, where x decreases every time the opponent makes a new proposal.

Definition 2. *Let $(\omega_1, \omega_2, \ldots \omega_K)$ be the list of all offers in Ω sorted in order of decreasing utility for agent a_1, and let π be the permutation of the integers 1 to K, such that $(\omega_{\pi(1)}, \omega_{\pi(2)}, \ldots \omega_{\pi(K)})$ is the list of all offers sorted in order of decreasing utility for agent a_2. Furthermore, for any integer $i \in \{1, 2, \ldots K\}$ we define $x_i := u_1(\omega_i)$, and $y_i := u_2(\omega_{\pi(i)})$, and we define b to be the smallest such integer for which $\Omega_1^{x_b} \cap \Omega_2^{y_b} \neq \emptyset$. Then we call $\Omega_1^{x_b} \cap \Omega_2^{y_b}$ the **balance set**, and we define*

$$x^\beta := min\{u_1(\omega) \mid \omega \in \Omega_1^{x_b} \cap \Omega_2^{y_b}\}$$

$$y^\beta := min\{u_2(\omega) \mid \omega \in \Omega_1^{x_b} \cap \Omega_2^{y_b}\}$$

*The values x^β and y^β are called the **balance values** of a_1 and a_2 respectively.*

The intuitive idea is that if a_1 and a_2 both apply the MiCRO strategy, then x_i and y_i are the minimum utility values they are respectively willing to accept after they have both made $i - 1$ proposals, and $\Omega_1^{x_i}$ and $\Omega_2^{y_i}$ represent the sets of offers they are then respectively willing to accept. Initially, these sets will be disjoint, but as i increases, x_i and y_i decrease, which means that $\Omega_1^{x_i}$ and $\Omega_2^{y_i}$ grow larger. Then, in the b-th round their intersection becomes nonempty, so at that point there are some offers they are both willing to accept. The balance set is by definition the set of these offers, and so the utility values they receive from the accepted offer must be greater than or equal to their balance values.

We have argued in Section 2.2 that a good negotiator should concede towards a certain threshold value, but no further. We will now show that the agents' balance values act as the threshold values for MiCRO. Specifically, Proposition 1 shows that, against a *consistent* opponent, MiCRO never makes any agreements below its balance value[3], while Proposition 2 shows that, if necessary, MiCRO does concede all the way to it. But before we can state these propositions, we first need to define what it means to be 'consistent'.

Let us first explain this with an example. Suppose a seller initially asks a price of \$100, while a buyer offers only \$75. Then, the seller decides to drop his price and ask \$50. Clearly, this would be silly, since the buyer has already indicated she is willing to pay \$75. And even if the buyer's offer of \$75 is no longer valid, the seller should at least try to re-propose the offer of \$75 first, before dropping to \$50. We therefore say the seller is making an *inconsistent proposal*.

Definition 3. *We say an agent a_i makes an* ***inconsistent proposal*** *if it proposes an offer ω after rejecting a better or equal offer ω' (i.e. $u_i(\omega) \leq u_i(\omega')$), unless a_i itself has already re-proposed ω'. Similarly, we say an agent a_i makes an* ***inconsistent acceptance*** *if it accepts an offer ω after rejecting a better offer ω', unless a_i itself has already re-proposed ω'. We say an agent a_i makes an* ***inconsistent rejection*** *whenever it rejects an offer ω' that is better than or equal to some offer ω that a_i itself has already proposed earlier. Finally, we say an agent is* ***consistent*** *if it never makes any inconsistent proposals, acceptances, or rejections.*

We argue that it is reasonable to assume that most agents act consistently most of the time, although we are certainly not claiming that this is *always* the case.

Proposition 1. *If an agent a_1 applies MiCRO and its opponent is consistent, then they will never make any agreement ω for which $u_1(\omega)$ is below a_1's balance value.* □

Proposition 2. *Let d be any bilateral negotiation domain in which the agents' balance values lie above their respective reservation values. Then, for at least one of the two utility functions of that domain, if MiCRO has that utility function, then there exists a strategy for the opponent that will force MiCRO to concede all the way to its balance value.* □

Now that we have established that the balance values act as threshold values for MiCRO, the question is to what extent these thresholds are optimal. A commonly used definition for the optimal solution in a negotiation is the Nash bargaining solution (NBS) [Nash, 1950]. Surprisingly, it turns out that the balance values of the ANAC domains indeed often lie very close to the NBS, even though we do not see any obvious reason why this should be the case. To quantify this, we can define a similarity score $s := \max\{|x^N - x^\beta|, |y^N - y^\beta|\}$ where x^N and y^N denote the utility values of the NBS. The lower this score, the closer the balance values are to the Nash values. If we calculate it for the 35 domains[4] used in ANAC 2012 and 2013 then we find that for 29 of them we have $s \leq 0.1$ and for 20 of them we even have $s \leq 0.05$ (in these domains for each agent the worst possible offer always has utility value 0.0 and the best possible offer always has utility value 1.0).

The following proposition may give some intuition as to why this peculiar fact holds. Specifically, it tells us that it is related to a kind of symmetry of the domains.

Proposition 3. *Let ω^* be some offer with utility values x^*, y^* (i.e. $x^* := u_1(\omega^*)$ and $y^* := u_2(\omega^*)$). Then, if ω^* is Pareto-optimal and $|\Omega_1^{x^*}| = |\Omega_2^{y^*}|$ then x^* and y^* are exactly the balance values.* □

Note that if ω^* is the NBS (or any other solution that one could consider as being 'optimal') then it is indeed Pareto-optimal. So, if the domain happens to satisfy the symmetry condition $|\Omega_1^{x^*}| = |\Omega_2^{y^*}|$, then the utility values of this optimal solution coincide exactly with the balance values.

6 Conclusions

We have presented a very simple new concession strategy, called MiCRO, which does not require any form of opponent modeling. We have shown that, despite its simplicity, it clearly outperforms all the top agents of ANAC 2012, 2013, 2018 and 2019. We therefore argue that the linear ANAC domains are too simplistic to truly assess the strength of complex negotiation algorithms. Our theoretical analysis shows that they can be tackled so easily because they are, in a certain sense, symmetrical with respect to the NBS. This knowledge may help researchers to design more challenging negotiation test cases in the future.

In summary, we draw two main conclusions:

1. New negotiation algorithms should preferably be tested on domains that are more challenging than the linear ANAC domains.
2. If one still insists on testing a new negotiation algorithm using linear ANAC domains, then at least the algorithm should be compared to MiCRO.

Furthermore, we argue that one may need to reassess the importance of existing opponent modeling algorithms if they have only been tested on linear ANAC domains.

Finally, we feel we should stress that we did not intend MiCRO to be a strategy for practical applications or for negotiations with humans. Instead, its main purpose is to serve as a benchmark strategy to assess the strength of other strategies.

Acknowledgments

This work was supported by a Juan de la Cierva - Incorporación research grant from the Spanish Ministry of Science and Innovation (IJC2018-036443-I).

[3] MiCRO does sometimes make proposals below its balance value, but Proposition 1 shows that a consistent opponent would never accept those.

[4] With their reservation values set to 0.

References

[Aydogan *et al.*, 2020] Reyhan Aydogan, Tim Baarslag, Katsuhide Fujita, Johnathan Mell, Jonathan Gratch, Dave de Jonge, Yasser Mohammad, Shinji Nakadai, Satoshi Morinaga, Hirotaka Osawa, Claus Aranha, and Catholijn Jonker. Research challenges for the automated negotiating agents competition (anac) 2019. In *Multi-Agent Systems and Agreement Technologies. 17th International Conference EUMAS 2020 and 7th International Conference AT 2020. Thessaloniki, Greece September 14-15, 2020. Revised Selected Papers*. Springer, 2020.

[Baarslag *et al.*, 2012] Tim Baarslag, Koen V. Hindriks, Catholijn M. Jonker, Sarit Kraus, and Raz Lin. The first automated negotiating agents competition (ANAC 2010). In *New Trends in Agent-Based Complex Automated Negotiations*, volume 383 of *Studies in Computational Intelligence*, pages 113–135. Springer, 2012.

[Baarslag *et al.*, 2013] Tim Baarslag, Katsuhide Fujita, Enrico H. Gerding, Koen V. Hindriks, Takayuki Ito, Nicholas R. Jennings, Catholijn M. Jonker, Sarit Kraus, Raz Lin, Valentin Robu, and Colin R. Williams. Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. *Artif. Intell.*, 198:73–103, 2013.

[Baarslag *et al.*, 2016] Tim Baarslag, Mark J. C. Hendrikx, Koen V. Hindriks, and Catholijn M. Jonker. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Auton. Agents Multi Agent Syst.*, 30(5):849–898, 2016.

[Bakker *et al.*, 2019] Jasper Bakker, Aron Hammond, Daan Bloembergen, and Tim Baarslag. RLBOA: A modular reinforcement learning framework for autonomous negotiating agents. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*, pages 260–268, 2019.

[Chen and Weiss, 2013] Siqi Chen and Gerhard Weiss. An efficient automated negotiation strategy for complex environments. *Engineering Applications of Artificial Intelligence*, 26(10):2613–2623, 2013.

[Faratin *et al.*, 1998] Peyman Faratin, Carles Sierra, and Nicholas R. Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3-4):159 – 182, 1998. Multi-Agent Rationality.

[Fujita *et al.*, 2014] Katsuhide Fujita, Reyhan Aydogan, Tim Baarslag, Takayuki Ito, and Catholijn M. Jonker. The fifth automated negotiating agents competition (ANAC 2014). In *Recent Advances in Agent-based Complex Automated Negotiation [revised and extended papers from the 7th International Workshop on Agent-based Complex Automated Negotiation, ACAN 2014, Paris, France, May 2014]*, volume 638 of *Studies in Computational Intelligence*, pages 211–224. Springer, 2014.

[Fujita *et al.*, 2017] Katsuhide Fujita, Reyhan Aydoğan, Tim Baarslag, Koen Hindriks, Takayuki Ito, and Catholijn Jonker. The sixth automated negotiating agents competition (anac 2015). In *Modern Approaches to Agent-based Complex Automated Negotiation*, pages 139–151. Springer, 2017.

[Gal and Ilany, 2015] Ya'akov (Kobi) Gal and Litan Ilany. The fourth automated negotiation competition. In *Next Frontier in Agent-Based Complex Automated Negotiation*, volume 596 of *Studies in Computational Intelligence*, pages 129–136. Springer, 2015.

[Jonge *et al.*, 2019] Dave de Jonge, Tim Baarslag, Reyhan Aydoğan, Catholijn Jonker, Katsuhide Fujita, and Takayuki Ito. The challenge of negotiation in the game of diplomacy. In *Agreement Technologies, 6th International Conference, AT 2018, Bergen, Norway, December 6-7, 2018, Revised Selected Papers*, volume 11327 of *Lecture Notes in Computer Science*, pages 100–114, Cham, 2019. Springer International Publishing.

[Lin *et al.*, 2014] Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen Hindriks, and Catholijn M. Jonker. Genius: An integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence*, 30(1):48–70, 2014.

[McKelvey *et al.*, 2017] Richard D. McKelvey, Andrew M. McLennan, and Theodore L. Turocy. Gambit: Software tools for game theory, version 15.1.1, 2017.

[Mirzayi *et al.*, 2021] Sahar Mirzayi, Fattaneh Taghiyareh, and Faria Nassiri-Mofakham. An opponent-adaptive strategy to increase utility and fairness in agents' negotiation. *Applied Intelligence*, 2021.

[Nash, 1950] J.F. Nash. The bargaining problem. *"Econometrica"*, "18":155–162, 1950.

[Sengupta *et al.*, 2021] Ayan Sengupta, Yasser Mohammad, and Shinji Nakadai. An autonomous negotiating agent framework with reinforcement learning based strategies and adaptive strategy switching mechanism. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '21, Online, May 5-7, 2021*. International Foundation for Autonomous Agents and Multiagent Systems, 2021.

[Williams *et al.*, 2011] Colin R. Williams, Valentin Robu, Enrico H. Gerding, and Nicholas R. Jennings. Using gaussian processes to optimise concession in complex negotiations against unknown opponents. In *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 432–438. IJCAI/AAAI, 2011.

[Williams *et al.*, 2014] Colin R. Williams, Valentin Robu, Enrico H. Gerding, and Nicholas R. Jennings. An overview of the results and insights from the third automated negotiating agents competition (ANAC2012). In *Novel Insights in Agent-based Complex Automated Negotiation*, volume 535 of *Studies in Computational Intelligence*, pages 151–162. Springer, 2014.